A Timeline of Model Development in ff15ipq

Karl T. Debiec∗, David S. Cerutti

April 3, 2016

Department of Chemistry, Michigan State University 578 South Shaw Ln., East lansing, MI 48823

∗ To whom correspondence should be addressed.

1 Overview

If pressed to give details, most force field developers will admit that their methods are somewhat more elaborate than the concise descriptions offered when the parameters are finally published. Occasionally investigators will describe frustration with the pace of progress, which can span years devoted by multiple collaborators leading to merely incremental improvements. Every force field project involves some degree of trial and error, and ff15ipq is no exception: the success of ff15ipq is, instead, in the way that some of our best results were obtained by pushing the proposed methods to their logical conclusion. This supporting document is intended as a means of disclosure. In addition to this narrative, Figure 1 shows the major steps and dates spanning the summer of 2015 in which ff15ipq took shape, and we include several other figures and tables describing the validation of each candidate.

The development of ff15ipq began with the intention to make minor modifications in the ff14ipq force field[1]: nothing more than a parameter fit for the nonstandard amino acid norleucine. It was the discovery that ff14ipq drastically overestimated salt bridge propensities in the canonical amino acids that led to a more vigorous renovation of the model focusing on new polar hydrogen radii. Even then, the presumption was that altering the polar hydrogen parameters and including perhaps 20,000 conformations of the various amino acids would supplement the 65,000 conformations in the original training set to finish the work. By late May 2015, KT Debiec had accumulated this data with spare compute cycles while focusing on other components of his graduate research. When these new parameters were finally synthesized into a working model, however, we observed the α-helical K19 to unfold quickly (see the following section, "Earlier Candidates for ff15ipq"). In contrast, the β-hairpin GB1 appeared to be excessively stable (data not shown), much as it seemed to have been in ff14ipq.

More had to be done, and the first resort was to take a look at the angle flexing terms to be sure that they were truly in step with the rest of the MP2/cc-pVTZ energy surface. DS Cerutti had developed a means for fitting both angle equilibria and stiffness constants based on the work of Hopkins and Roitberg[2], but after seeing the recently published work of Kenno Vannomeslaeghe and Alex Mackerell[3], agreed that the CHARMM team had a more elegant implementation and added that to the mdgx package. When making a new OPLS protein model candidate, optimizing the N-C_α-C, C_α-C-N, and C-N-C_α angles had made the difference between a model that worked quite well and one that hardly worked at all. These three angles along the backbone were reoptimized as a trial run in ff15ipq, and the results were just as dramatic: agreement in NMR *J*-couplings

improved considerably, as shown in Table 1, α-helicies began to show stability, and behavior in the TrpCage miniprotein improved as well. In terms of sheer data fitting, optimizing these three angles made the biggest improvement in the ff15ipq candidate's ability to model the energy surfaces of small amino acids, as shown in Table 2.

However, the stability in many systems was not enough, and the β-hairpin appeared to retain too much stability at temperatures when it should have melted. We turned next to our aggressive increase of the polar hydrogen radius, from 1.06Åin the original Cornell model line to 1.5Å, and rolled this back to 1.3Å. This yielded further improvements, but the reoptimized values of the backbone angles, all considerably smaller than their initial values, were encouraging a net contraction of the backbone, puckering in the planar peptide groups, and very likely favoring some secondary structures over others. We wanted to improve performance while retaining naturally expected results such as that the three angles surrounding the peptide backbone N and C sum to 360°. Following the improvements of the first round of angle fitting, KT Debiec had wanted to push the optimization further, but DS Cerutti had pushed back, not wanting to saddle the optimization protocol with too many degrees of freedom. Debiec had been right: adding new conformations and including the backbone C_α-N-H and C_α-C=O angles in the bonded parameter fit led to further improvements and gave us a model that the final version of ff15ipq only rivals in performance. It was lingering concerns about the planarity of the peptide bond that led us to fully optimize angles along the backbone, including those involving hydrogen atoms and all six surrounding the C_α.

To be circumspect, the simultaneous optimization of all these angles did not *initially* yield a good model, as found by some short MD simulations which are not included here: another generation of conformations, optimized by that candidate, however, did. It is likely that including all of the angles together, three around the N and C atoms and six around C_α, introduced a tug-of-war which was not properly balanced by fitting the energy surfaces sampled along each parameter individually. Again, by letting the model drive and explore conformations that the combinations of new angles favored, we obtained enough results to improve the balance. As shown in Table 3, the final optimized values of all these backbone angles came very close to the initial sums set deliberately in the Amber force field even before the advent of the Cornell model line.

In retrospect, ff14ipq had one fatal flaw: the imbalance in hydration versus protein:protein interactions introduced by making polar heavy atoms interact differently with one another than they did with water. The rationale had been to simultaneously achieve the right hydration free energies without destroying our ability to fit torsion parameters, but this was the root of the aberrant salt bridge propensities and even a source of excess

stability in native secondary structures: we had unwittingly glued the model together, a fact which could have cast doubt on whether the IPolQ scheme was worthwhile at all. ff15ipq seeks to capture the same physics but without cutting any corners, the pursuit of natural backbone angle sums being another part of the effort. The result is satisfying in its own right: over 200μs of trajectories it seems to meet with the state of the art. While the bar for a new protein force field is very high, the fact that we delved into trial and error only to realize that the best choice was to push our proposed methods to their limits, exposing all of the relevant parameters to our optimization scheme, sends an inspiring message that the reasoned pursuit of physics did not fail.

2 On the Importance of Accurate Data Fitting

The design of ff15ipq rests heavily on data fitting, and each step of the evolution described in Figure 1 was accompanied by a large influx of new data. The final, published version of ff14ipq included some 65,000 amino acid conformations in its torsion parameter fitting data, and the first version of ff15ipq included more than twice as many. The reduction in accuracy for Glycine tetrapeptide shown in Table 2 reflects this tremendous growth in the training set, but overall the growth in ff15ipq torsion parameter space was sufficient to cover a larger conformational space while maintaining agreement over the limited domain used to train ff14ipq. The introduction of angle fitting shifted the balance. The new conformations added to sample angle flexion could be highly strained–departures of 10°from the original equilibria were common–but the model was able to learn better values for these degrees of freedom and improve the energy surfaces of many amino acids.

It is commonly assumed that, because bond stretching is governed by forces roughly ten times stiffer than angle flexion, the bonds can be fitted independently and taken as given when optimizing other parameters. This is probably true, not least because motions along bonds are orthogonal to motions along any of ther degree of freedom: with the exception of rings, atoms in a structure can be moved so as to strech bonds without changing any angles or torsions. A similar assumption is made for the angle flexing terms themselves, but while they are roughly an order of magnitude stiffer than individual torsion terms they are not necessarily that much higher in energy than the *sums* of all torsions about a particular bond, nor are they so neatly orthogonal to the torsional degrees of freedom. A prime example is the C_β–C_α–H_α angle: flexion in this angle projects heavily onto the backbone ϕ and ψ dihedral angles, which are among the most thoroughly refined parameters in modern force fields.

It is likely that including backbone angles in the data fitting helped keep the torsions from being saddled

with details of the energy surface that they were not well suited to address. The effects of this move are stunning in the reproduction of Ala(5) *J*-couplings, and also in the stability of alpha-helical peptides, despite the fact that our final round of refinement took us from an excellent result in the cases of K19 and $AAQAA_3$ to one with slightly less stability than would be ideal. The balance of α-helical and β-sheet stabilities will remain a difficult problem, but we are exploring ways of improving it, even in the context of nuclear-centered, non-polarizable charge models with the present descriptions of bonded interactions.

Throughout the development, there was a spirited debate between DS Cerutti and KT Debiec over the inclusion of parameters which appeared to have small effects on the overall accuracy of the fit. ff15ipq spans 750 torsion parameters due to further specialization of C_α atom types. One attempt at reducing the numbers of torsions describing χ_1 rotations and planarity in imidazole rings of His and Trp residues, not presented among the major developments, was quickly shown to yield a poorer result. One could argue that additional generations of refinement could have improved the outcome, a possibility we did not test. However, one can also take a different view of the error reported in Table 2. The average error in reproducing the gas-phase MP2 energy surface is a sum of many factors, but they can be roughly grouped into bonded and non-bonded interactions. Of the bonded interactions, stretching terms in the molecular mechanics potential have real disagreements with quantum physics, but as shown in earlier work[1] these disagreements do not have a hard impact on the *relative* energies. Of the non-bonded interactions, steric repulsion of atoms at short range is handled by a combination of 1:4 interactions, torsions, and angle flexion terms, while the truly long-ranged interactions are dominated by electrostatics. As has been documented[4], the errors in the electrostatic potential made by a non-polarizable, nuclear charge model can be quite significant. The additional angle flexing optimizations should be viewed in this context: bonded and non-bonded interactions constitute two orthogonal sources of the errors reported in Table 2, two sides of a right triangle for which we can only directly measure the hypotenuse. To some, a 10-15% improvement in the overall error may seem unexciting, but it could indicate a much larger improvement in the subset of bonded interactions. In future additions to the ff-ipq lineage, we intend to explore the advantages of non-nuclear charge centers that could dramatically enhance the electrostatic description and clarify the portions of the energy surface that the angle and torsion terms need to pick up.

3 Transitional Forms: The Steps Between Each Milestone

Once all of the machinery was in place, generating new data and fitting new models was relatively easy. Throughout the summer of 2015 new parameter sets were being created every seven to ten days, the four highlighted in Figure 1 being the most significant milestones. The validation efforts evolved over this time as well: new systems were introduced, beyond the original ff14ipq validation set, hydrogen mass repartitioning doubled our simulation timescales, and our concept of convergent results changed as well. The first few candidates' validation runs lasted on the order of 100ns, at which point most of the results could be found inaccurate. By the time ff15ipq V3 was minted, however, it had become necessary to simulate on the microsecond timescale to see whether qualitative agreements with experiment would grow into quantitative agreements and distinguish one candidate over another. We set out to make a force field appropriate for the microsecond timescale, and this is what we achieved. Here we offer more details on some of the candidates we passed up.

The V1 candidate was itself the product of meticulous efforts. Between the publication of ff14ipq and June 2015, KT Debiec had worked on ff15ipq as one of several graduate research projects and set up an extensive system for generating and archiving the growing data set, infrastructure that proved valuable as development continued. The low α-helical stability displayed by this candidate shown in Figure 2 led to rejection, even though we had observed reasonable TrpCage stability the simulations were relatively short.

With the introduction of angle fitting and subsequent reduction of the polar hydrogen radius back to 1.3Å, version 2 displayed slightly improved α-helical stability, but there still did not appear to be enough and the major backbone angles had shifted their equilibria significantly as shown in Table 3. The TIP4P-Ew water model[5], despite its excellent bulk properties, might be imbalanced for nuclear-centered protein charge models due to its reliance on a massless site for its negative charge and the subsequent increase in overall charges on that massless site and hydrogen atoms (most non-polarizable water models tend to have dipoles of roughly 2.35, but the TIP4P geometry implies that this dipole is constructed with charges that are approximately 20% stronger than the equivalent three-point water models). Hypothesizing that three-point water models may be better suited, we switched to SPC/Eb[6] form. This improved performance on GPU codes which do not make accommodations for atoms with only charge or steric properties, but it had no discernable effect on the stability of the K19 α-helix as shown in Figure 2. This is one indication that angle fitting was indeed a benefit to secondary structure stability. With only one angle around each backbone N, C_α, or C atom, the possible optimizations are either that the sum of angles depart from the preset circular or tetrahedral values or that only the stiffness constant of the

angle changes (the latter is not nearly as effective for improving the quality of data fitting). Around the nitrogen, some tendency away from planarity can be tolerated (the peptide hydrogen is known to lie slightly out of the plane of the bond),[7,8] but the carbonyl carbon does not offer much leeway. Seeking to restore the original angle sums about each of the backbone heavy atoms, we introduced two more angle optimizations about the peptide N and C atoms. The results marked the third milestone in ff15ipq development. As shown in Figures 2 and 3, this model displayed excellent α-helical stability, but on the microsecond timescale showed uncertain stability with GB1 (Figure 4). It is not clear whether GB1 was properly metastable in this model, but it was clear that the angle sums were still not brought back to the desired values (some difficulty in stabilizing this anti-parallel β-sheet may have stemmed from the angles imparting a net curl to the backbone). At the same time, a Chignolin variant was found to be completely stable even at very high temperatures (chignolin is itself a β-hairpin, but a much smaller one that would not be so susceptible to backbone curl).

The release version of ff15ipq met our goal of restoring the angle sums around the carbonyl carbon, and slightly undershot the sum for the backbone nitrogen to impart a subtle tendency away from planarity. It is notable that simultaneous optimization of all the angles around the peptide bond produced results much closer to the naturally expected values than limiting the optimization to just two of the three angles around each heavy atom. It is also notable that the high quality of our Ala(5) results in Table 1 was unaffected–suggesting that, while the angle values in the earlier Cornell model line tended to produce only moderately accurate results, it is not necessary to distort the overall character of the backbone to improve the situation either. Although the α-helical stability in the K19 and $AAQAA_3$ systems appears diminished relative to the V3 candidate, the simulations in Figure 3 are only 2μs whereas the simulations of the V4 candidate stretch twice as long: it is actually ambiguous whether one candidate is better than the other at depicting α-helical stability. V4 does show refolding in the GB1 Hairpin system whereas V3's shorter simulations do not. As shown in the main text, V4 may show a slight temperature dependence in the stability of the Cln025 Chignolin mutant, but this observation again rests on significantly longer simulation lengths than were attempted with V3. V4 does show refolding in the GB1 system whereas V3's shorter simulations do not. The overall quality of results and completeness of parameter optimization–all torsions, and every angle along the backbone–in V4 led us to choose this candidate.

References

[1] Cerutti, D.; Swope, W.; Rice, J.; Case, D. *J. Chem. Theory Comput.* **2014**, *10*, 4515–4534.

[2] Hopkins, C.; Roitberg, A. *J. Chem. Inf. Model.* **2014**, *54*, 1978–1986.

[3] Vannomeslaeghe, K.; Yang, M.; Jr., M. A. *J. Comput. Chem.* **2015**, *36*, 1083–1101.

[4] Cerutti, D.; Swope, W.; Rice, J.; Case, D. *J. Phys. Chem. B* **2013**, *117*, 2328–2338.

[5] Horn, H.; Swope, W.; Pitera, J.; Madura, J.; Dick, T.; Hura, G.; Head-Gordon, T. *J. Chem. Phys.* **2004**, *120*, 9665–9678.

[6] Takemura, K.; Kitao, A. *J. Phys. Chem. B* **2012**, *116*, 6279–6287.

[7] Hu, J.-S.; Bax, A. *J. Am. Chem. Soc.* **1997**, *119*, 6360–6368.

[8] Mannfors, B.; Mirkin, N.; Palmo, K.; Krimm, S. *J. Phys. Chem. A* **2003**, *107*, 1825–1832.

[9] Graf, J.; Nguyen, P.; Stock, G.; Schwalbe, H. *J. Am. Chem. Soc.* **2007**, *129*, 1179–1189.

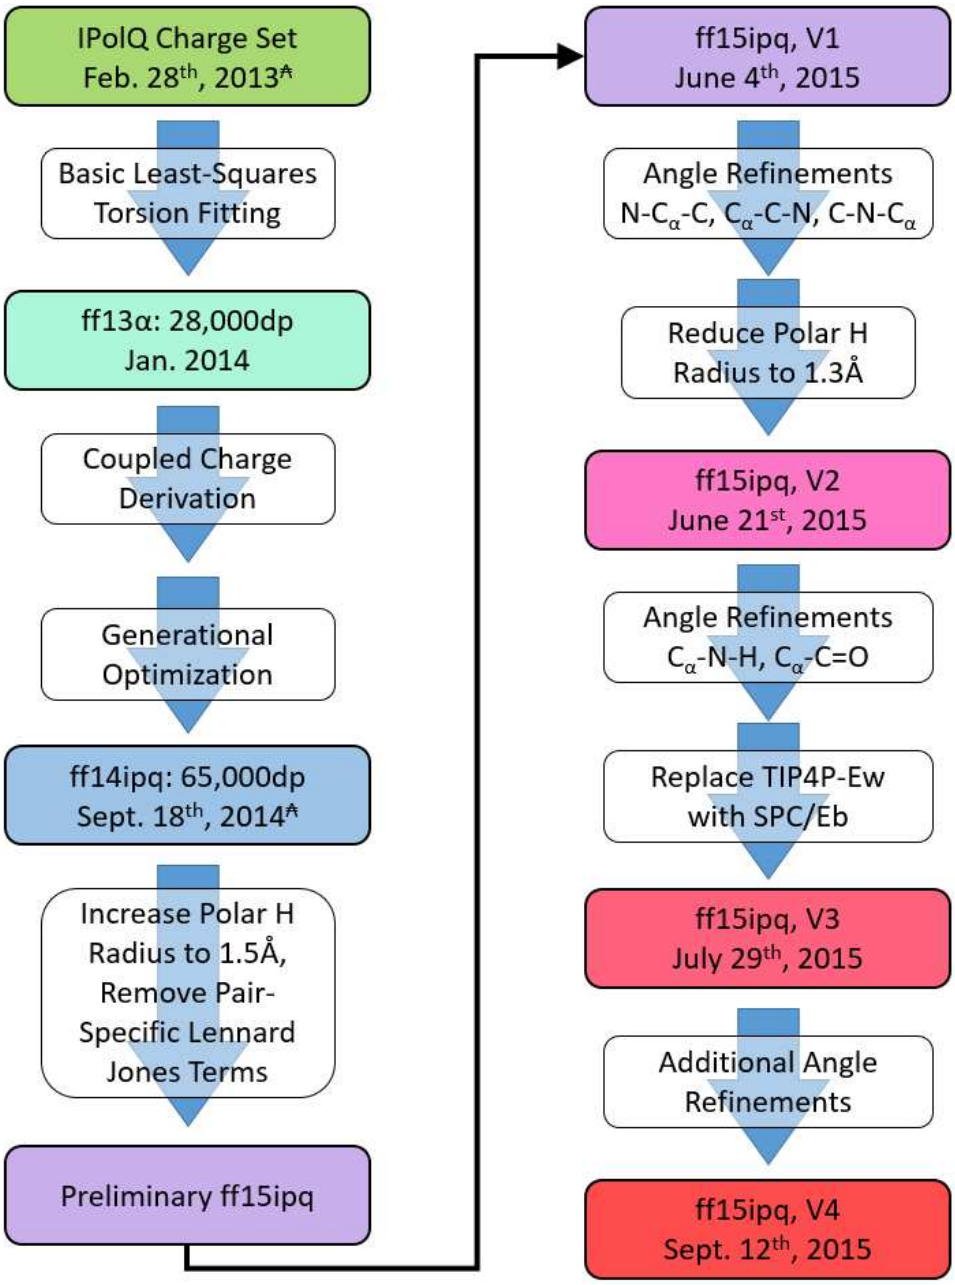

Figure 1: **A pictorial history of ff15ipq.** While ff15ipq has unique parameters in almost every meaningful respect, it attempts to capture the same physics as the earlier ff14ipq. Carat symbols indicate dates of publications; other dates are taken from time stamps on parameter files and outputs of mdgx. (The mdgx outputs for the final version are included with the Supporting Information.) With each change in polar hydrogen radius, over 750 partial charges for the canonical amino acids (plus their protonated and terminal forms) were reoptimized. Changes in the radii as well as inclusion of new optimizable angles triggered reoptimization of nearly 900 torsion parameters (as well as any angle parameters previously subject to fitting).

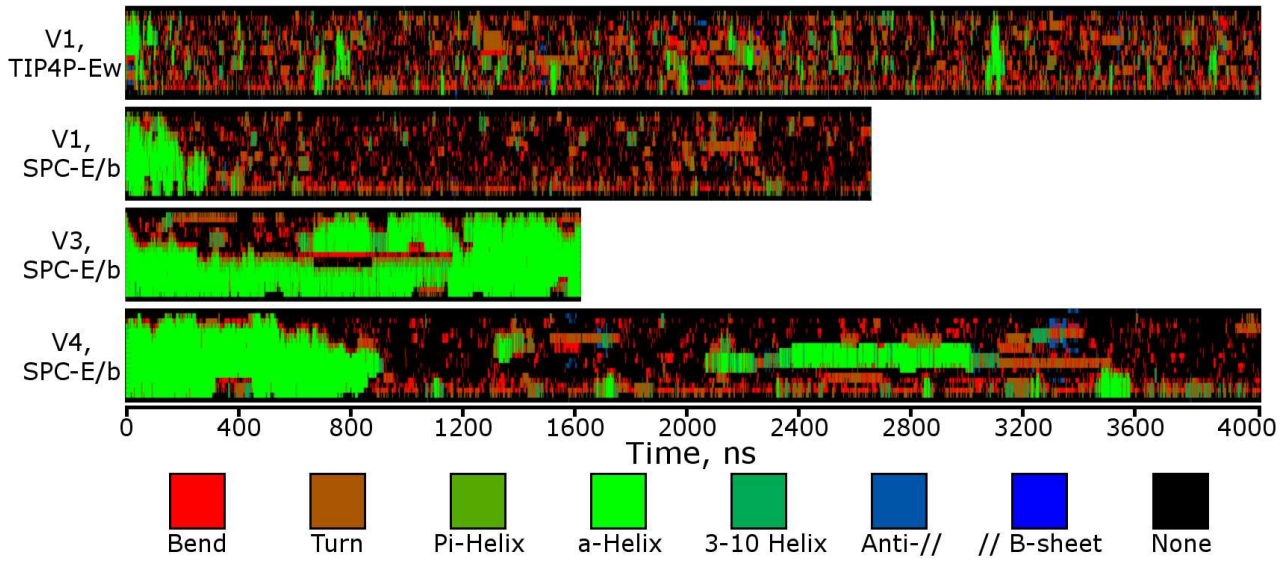

Figure 2: **Simulations of the K19 System with Several Candidates.** All simulations shown in this plot were carried out at 275K. The K19 peptide is Ace-GGGKAAAAKAAAAKAAAAK-Nhe, although the first three candidates simulated a system with Nme, not a bare amide group at the tail (we found the effect to be minimal in subsequent simulations with the release version). Secondary structures of each residue in the sequence was calculated by DSSP and plotted in ascending order on the *y*-axis, evolving with time along the *x*-axis. The first candidate, lacking any angle optimization, did not stabilize the peptide (our results suggest that the only way in which the very similar ff14ipq was able to stabilize this peptide was the "stickiness" discussed in its initial publication, which strengthened hydrogen bonding between backbones just as much as between backbones and side chains). In this case, changing the water model does not appear to have helped the stability, but the introduction of more extensive backbone angle optimization in the third candidate did improve the situation. Some of this progress had to be sacrificed in the release version, but the helicity has not completely vanished, suggesting that we are near the correct helical stability.

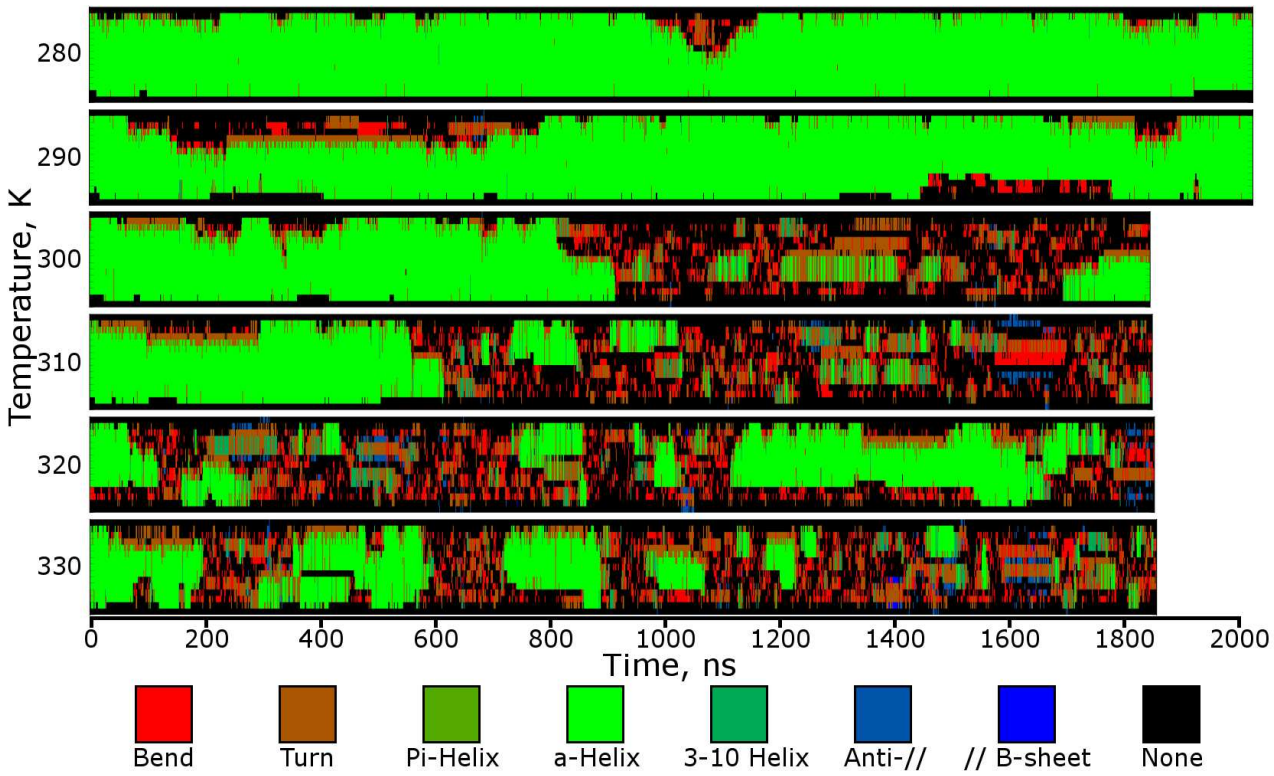

Figure 3: **Simulations of the AAQAA$_3$ System with the V3 Candidate.** For each temperature, the residues in the sequence Ace-AAAQAAAAQAAAAQAA-Nme are plotted in ascending order on the *y*-axis, with the backbone secondary structure classification color coded as the system evolved with time. Over simulations of 1μs length, this α-helical system exhibits high stability at 280K, and descreasing stability as the temperature increases. Cooperative formation and dissolution of the helix is observed multiple times at temperatures greater than 300K. In the interest of full disclosure, this system contained an extra Ala residue just after the acetyl blocking group at its N-terminus, and a methylamide blocking group at the end rather than the experimental system's amide. These errors in system preparation were corrected before final validation on the V4 candidate, but we did not observe any significant conseqeunces on α-helical integrity due to these changes.

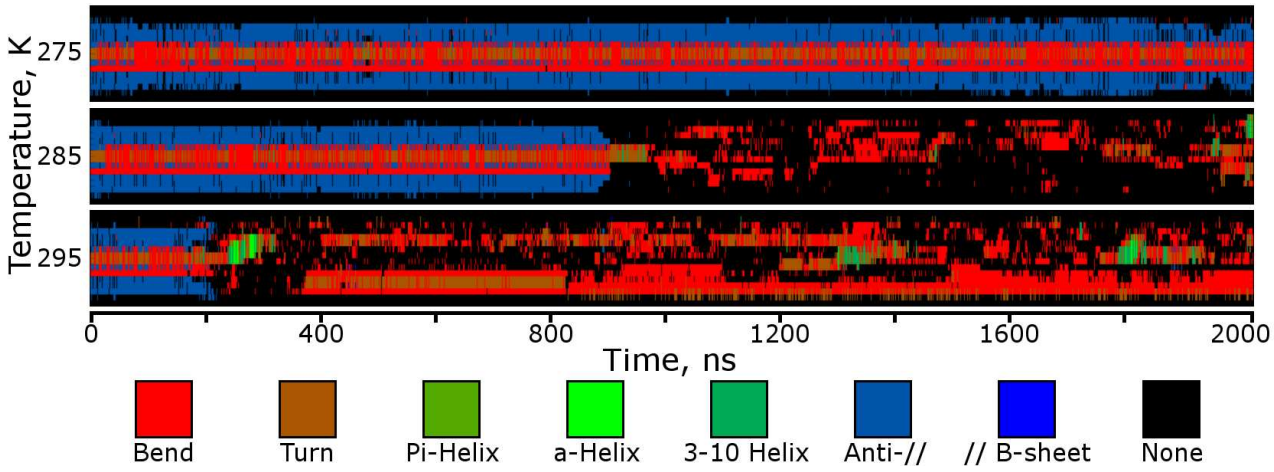

Figure 4: **Simulations of the GB1 Hairpin with the V3 Candidate.** For each temperature, the residues in the sequence Ace-GEWTYDDATKTFPTVTE-Nme are plotted in ascending order on the *y*-axis, with the backbone secondary structure classification color coded as the system evolved with time. Our 2μs simulations show this β-sheet system eventually unfolding but not (yet) refolding.

J-Coupling	Residue	Experimental	ff15ipq Model χ^2			
		Value	1[a]	2	3	4
^1J(N, C_α)	2	11.36 ± 0.59	0.23 ± 0.03	0.02 ± 0.01	0.00 ± 0.00	0.00 ± 0.00
^1J(N, C_α)	3	11.26 ± 0.59	0.70 ± 0.17	0.47 ± 0.13	0.28 ± 0.06	0.11 ± 0.01
^2J(N, C_α)	2	9.20 ± 0.50	8.92 ± 0.15	1.11 ± 0.01	1.16 ± 0.01	1.24 ± 0.00
^2J(N, C_α)	3	8.55 ± 0.50	0.36 ± 0.07	0.06 ± 0.02	0.04 ± 0.01	0.00 ± 0.00
^3J(C , C)	2	0.19 ± 0.22	8.21 ± 0.36	5.89 ± 0.27	4.52 ± 0.08	4.69 ± 0.12
^3J(H_α, C)	2	1.85 ± 0.38	0.17 ± 0.09	0.03 ± 0.06	0.33 ± 0.23	0.53 ± 0.17
^3J(H_α, C)	3	1.86 ± 0.38	0.28 ± 1.17	0.35 ± 0.51	0.03 ± 0.30	0.01 ± 0.01
^3J(H_N, C)	2	1.10 ± 0.59	0.07 ± 0.01	0.04 ± 0.01	0.06 ± 0.01	0.08 ± 0.01
^3J(H_N, C)	3	1.15 ± 0.59	0.01 ± 0.01	0.00 ± 0.01	0.00 ± 0.00	0.01 ± 0.00
^3J(H_N, C_β)	2	2.30 ± 0.39	1.11 ± 0.09	0.70 ± 0.08	0.30 ± 0.03	0.26 ± 0.02
^3J(H_N, C_β)	3	2.24 ± 0.39	1.12 ± 0.14	0.91 ± 0.10	0.48 ± 0.10	0.41 ± 0.04
^3J(H_N, H_α)	2	5.59 ± 0.91	0.07 ± 0.02	0.02 ± 0.01	0.03 ± 0.01	0.07 ± 0.01
^3J(H_N, H_α)	3	5.74 ± 0.91	0.17 ± 0.04	0.14 ± 0.04	0.01 ± 0.01	0.00 ± 0.00
^3J(H_N, C_α)	2	0.67 ± 0.10	0.30 ± 0.02	0.30 ± 0.04	0.45 ± 0.02	0.20 ± 0.01
^3J(H_N, C_α)	3	0.68 ± 0.10	0.42 ± 0.09	0.24 ± 0.04	0.36 ± 0.04	0.31 ± 0.02
Time, ns			1472	1222	2928	6000
Mean χ^2	ALL		1.47 ± 0.07	0.69 ± 0.04	0.54 ± 0.02	0.53 ± 0.02

Table 1: **NMR J-Couplings calculated from simulations of the Ala(5) system.** Results are shown for the original Karplus coefficients[9]. Data from a version of the force field created after angle fitting was introduced but before the polar hydrogen radius was reduced was lost, but the results were consistent with later versions. Angle fitting was the major driver of improvement in this test: the decrease in polar hydrogen radius affected the stability of larger secondary structures much more than conformations of this penta-peptide.

[a] Model indices are indicated in Figure 1

	Model Accuracy against Initial Data Set, kcal/mol Root Mean Squared Error					Final Data Set kcal/mol RMSE	
Amino Acid	ff14ipq	ff15ipq Candidate				ff14ipq	ff15ipq
		V1	V2	V3	V4		
Ala_3	1.17	1.16	1.02	0.97	1.03	2.29	1.04
Gly_3	0.96	1.07	0.95	0.94	0.97	2.58	0.86
Arg	1.09	1.04	1.08	1.07	1.17	>5.0	2.21
Asn	0.90	0.94	0.89	0.88	0.88	1.50	1.15
Asp	1.90	1.78	1.51	1.45	1.49	2.04	1.54
Cys	1.02	1.06	1.09	1.03	1.00	1.46	1.05
Gln	0.72	0.75	0.75	0.75	0.74	1.32	0.97
Glu	1.61	1.60	1.47	1.47	1.64	2.35	1.83
Hie	0.87	0.87	0.82	0.78	0.77	1.54	1.03
Leu	0.88	0.84	0.86	0.85	0.86	1.15	0.78
Phe	0.82	0.89	0.91	0.87	0.79	1.31	0.93
APA[a]	1.64	1.38	1.33	1.31	1.12	3.51	1.43
Ser	0.84	0.83	0.82	0.81	0.79	1.34	1.03
Val	0.70	0.61	0.62	0.67	0.59	1.27	0.77

Table 2: **Errors for fitting conformations of selected amino acids from the original ff14ipq data set and the final ff15ipq data set.** The initial data set (28,000 conformations) used in training ff14ipq was re-scored with evolving ff15ipq model candidates. All systems are blocked di- or tetra-peptides; for the purposes of this comparison to the relative single-point energies of vacuum-phase MP2/cc-pVTZ calculations, the vacuum-phase charge set, coupled with torsions trained in the context of those charges, was used. Results in the ff14ipq column differ from those reported earlier [1], as the earlier results concerned the first generation of ff14ipq against its own fitting data; the final model of ff14ipq, reported here, was fitted to a somewhat larger data set (65,000 conformations, comprising the original 28,000). The first ff15ipq candidate was trained on these and an additional 90,000 conformations; the final data set included over 263,000. The rightmost columns depict the accuracy of the published ff14ipq and ff15ipq (V4) force fields when challenged with the full 263,000 conformations.

[a] Ala-Pro-Ala blocked tetrapeptide, not included in the original ff14ipq training set. Optimizations of angles specific to proline were finally introduced in the most recent version of ff15ipq.

	Neutral, Non-Proline, Branched Residues		
Version	Amide N (Target 360)	C_α (Target 657)	Carbonyl C (Target 360)
V1	359.9	659.4	359.9
V2	357.5	655.4	356.8
V3	352.9	654.4	353.4
V4	356.4	667.0	359.6
	Glycine		
V1	359.9	659.4	359.9
V2	353.3	656.4	355.0
V3	347.6	656.1	354.1
V4	347.8	661.3	359.7
	Cationic Amino Acids		
V1	359.9	659.4	359.9
V2	356.4	655.1	359.0
V3	347.4	654.1	358.1
V4	345.1	663.1	362.4
	Anionic Amino Acids		
V1	359.9	659.4	359.9
V2	362.8	653.3	358.5
V3	357.2	652.4	348.8
V4	351.0	660.8	357.1

Table 3: **Sums of angles around specific backbone atoms.** The canonical peptide backbone structure has planar geometry around the amide nitrogen and carbonyl carbon atoms, and tetrahedral geometry around the core C_α atom. In the original Cornell force field and subsequent edits such as ff99SB and ff14SB, the angles around the peptide bond sum to 360°by construction. If the sum of three angle equilibria in a group is dramatically less than 360, there will be a tendency for the group to pucker out of plane (three angle equilibria of 109.5°would create a tetrahedral arrangment). Similarly, distortions of a tetrahedron can make the sum of six angles around it sum to less than 657°. Sums greater than the target values imply that there is pressure between angles at equilibrium, which helps to enforce planarity of the peptide bond and tetrahedral character around the C_α.